CHAPTER 4
THIRD GENERATION HF MESSAGING PROTOCOL

4.1 Introduction

4.1.1 Summary

The United States Department of Defense is currently in the process of updating U.S. MIL-STD-188-141A, the automated HF radio standard commonly referred to as the ALE standard. Ratification of the new version, to be known as MIL-STD-188-141B, is planned for late 1998. Appendix C of the new version defines a unified third generation synchronous HF messaging protocol (3G-HF) that uses burst serial tone waveforms for link establishment and link maintenance, as well as improved data link engines for the faster transfer of data traffic. The third generation unified messaging protocol, henceforth named third generation HF, is intended for HF networks with intense voice and/or data message loading.

This paper begins with an overview of the layered architecture. This is followed by a brief overview of the constituent waveforms within the physical layer and an overview of the data link layer (DLL). The data link layer overview includes descriptions of the connection set-up (CSU) and traffic set-up (TSU) protocols, the high-rate and low-rate data link protocols, and the circuit link protocol. Finally, the paper discusses a framework for simulation of the entire HF messaging system.

4.1.2 Background

The need for a next-generation HF messaging protocol suite stems from the growing need for HF voice and data messaging systems offering high reliability and high capacity. HF continues to play a crucial role in military beyond-line-of-sight communications; however, this role is changing as more HF military users extend the internet into the battlefield. The use of standard internet applications (such as E-mail) over wireless transmission media (specifically HF) creates heightened technical challenges which are not met adequately by existing HF communications protocols. The existing protocols do not provide effective channel access mechanisms, and, as a result, tend to break down due to collisions and congestion under heavy network loads. The current ALE and data link standards use very different modulation formats (8-ary FSK vs. serial tone PSK), resulting in a performance mismatch between the linking subsystem and the message delivery subsystem. Current HF ARQ protocols require complicated methods for matching the waveform and/or data rate to the channel conditions. The 3G-HF system attempts to meet all of these challenges with simple but effective designs for:

- prioritized channel access and collision avoidance,
- a unified and scaleable burst waveform design used for connection set-up, traffic set-up, and traffic exchange,
- and an enhanced ARQ protocol that significantly increases throughput in all channel conditions, while simplifying the data rate adaptation algorithm.

A key element of any successful and widely implemented protocol is simplicity. One only has to look to the internet protocols to see that simple protocols thrive in the marketplace at the expense of more complex alternatives. While any interoperability standard requires correctness, completeness, conciseness, consistency, and clarity, the goal of simplicity must not be sacrificed to achieve the other goals. This is a major guiding principle applied in the definition of the 3G-HF protocols.

Figure 4.1 shows the concept architecture for 3G-HF. A specific architecture implementation is not required by 3G-HF; nor is the implementation of each service primitive. The protocol architecture and service primitives are provided within the standard for completeness and to aid in specifying the required over-the-air behavior in the form of exchanges of protocol data units (PDUs). In 3G-HF an HF subnet may have a network layer (NL), a data link layer (DLL), and a physical layer (PL). This does not preclude the existence of transport, session, presentation, and even application layers within the system. The 3G-HF specification pertains only to the DLL and PL. The DLL consists of a connection manager (CM), a traffic manager (TM), data link protocols (DLPs), and a circuit link protocol (CLP). Overviews of the PL and of each component of the DLL are provided in subsequent sections.

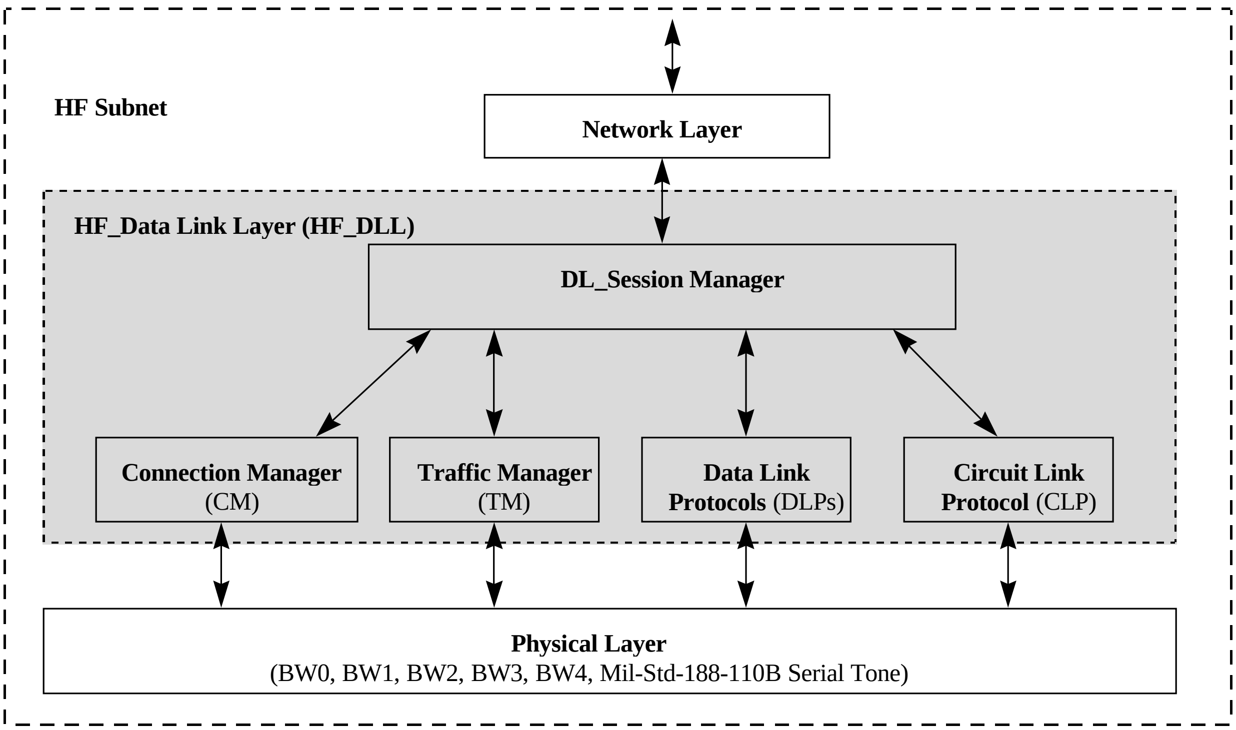

FIGURE 4.1

Conceptual view of the third generation architecture

It is important to note that Appendix C of MIL-STD-188-141B was still in preparation as this paper was being written; as a result, the protocol definitions and performance data presented here represent a ‘snapshot’ of 3G-HF taken at one moment in time, and may change by the time the standard is ratified. The paper is intended to provide a useful overview of the basic concepts and salient features of the 3G-HF protocols defined in Appendix C. Space limitations make it impossible to present many protocol details in this paper. These will, however, be presented in full in the published MIL-STD-188-141B specification.

4.2 Constituent waveforms

The 3G-HF system uses a family of scaleable burst waveform signaling formats for transmission of all control and data traffic signaling. Reference [1] provides a more detailed overview of these burst modems and their performance. Scaleable burst waveforms are defined for the various kinds of signaling required in the system, so as to meet their distinctive requirements as to payload, duration, time synchronization, and acquisition and demodulation performance in the presence of noise, fading, and multipath. All of the burst waveforms use the basic 8-ary PSK serial tone modulation at 2400 symbols per second that is also used in the MIL-STD-188-110A serial tone modem waveform. The low-level modulation and demodulation techniques required for the new system are similar to those of the 110A modems. TABLE 4.1 provides an overview of the characteristics of the various burst waveforms.

In contrast to the MIL-STD-188-110A waveform, the waveforms used in the 3G-HF system are designed to balance the potentially conflicting objectives of maximizing the time-diversity achieved through interleaving, and minimizing on-air time and link turn-around delay. The latter objective plays an important role in improving the performance of ALE and ARQ systems, which by their nature require a high level of agility.

TABLE 4.1

Waveform characteristics overview

Wave-form	used for	burst duration	pay load	preamble	FEC coding	inter-leaving	data format	effective code rate
BW0	Connection Set-Up (CSU) PDUs	613.33 ms 1472 PSK symbols	26 bits	160.00 ms 384 PSK symbols	rate = 1/2, k = 7 convolutional (no flush bits)	4x13 block	16-ary orthogonal Walsh function	1/96
BW1	Traffic Set-Up (TSU) PDUs; High-Rate Data Link acknow-ledgement PDUs	1.30667 sec 3136 PSK symbols	40 bits	240.00 ms 576 PSK symbols	rate = 1/3, k = 9 convolutional (8 flush bits)	16x9 block	16-ary orthogonal Walsh function	1/144
BW2	High-Rate Data Link traffic data PDUs	640 + (n*400) ms 1536 + (n*960) PSK symbols, n = 3, 6, 12, or 24	n*188 1 bits	26.67 ms 64 PSK symbols (for equalizer training)	rate = 1/4, k = 8 convolutional (7 flush bits)	none	32 unknown/ 16 known	variable: 1 / 1 to 1 / 4
BW3	Low-Rate Data Link traffic data PDUs	373.33 + (n*13.33) ms 32n + 896 PSK symbols, n = 64, 128, 256, or 512	8n+25 bits	266.67 ms 640 PSK symbols	rate = 1/2, k = 7 convolutional (7 flush bits)	24x24, 32x34, 44x48, or 64x65 convol-utional block	16-ary orthogonal Walsh function	variable: 1 / 12 to 1 / 24
BW4	Low-Rate Data Link acknow-ledgement PDUs	640.00 ms 1536 PSK symbols	2 bits	none	none	none	4-ary orthogonal Walsh function	1 / 1920

It is illuminating to compare the probability of linking under these conditions for the 3G-HF system vs. MIL-STD-188-141A ALE. A connection set-up (CSU) requires a two-way exchange of CSU PDUs conveyed using the BW0 waveform. As shown, the 3G-HF system is likely to achieve a probability of linking equivalent to that of 141A ALE, under conditions of roughly 5-7 dB lower signal-to-noise ratio (SNR).

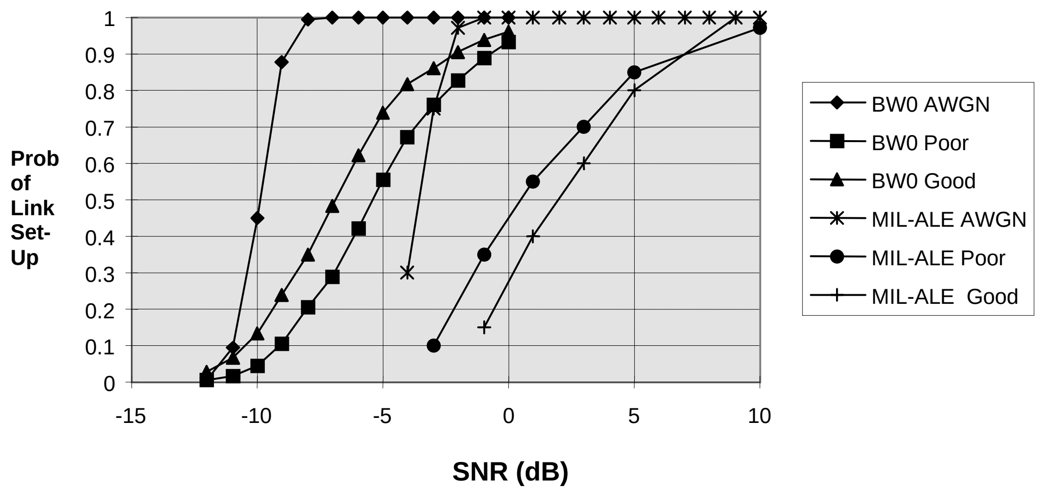

FIGURE 4.2

BW0 probability of linking vs. MIL-ALE

4.5 Connection manager

The connection manager is responsible for the connection set-up phase (otherwise known as ALE). Third-generation ALE (3G-ALE) is designed to quickly and efficiently establish one-to-one and one-to-many (both broadcast and multicast) links. It uses a specialized carrier-sense-multiple-access (CSMA) scheme to share *calling channels*, and monitors *traffic channels* prior to using them to avoid interference and collisions. Calling and traffic channels may share frequencies, but the system is likely to achieve better performance when they are separate. Each calling channel is assumed to be associated with one or more traffic channels that are sufficiently near in frequency to have similar propagation characteristics. The concept of associated control and traffic frequencies can be reduced to the case in which the control and traffic frequencies are identical.

4.6 Scanning

As in second-generation ALE (2G-ALE), 3G-HF receivers continuously scan an assigned list of calling channels, listening for 2G or 3G calls. However, 2G-ALE is an asynchronous system in the sense that a calling station makes no assumption about when a destination station will be listening to any particular channel. The 3G-HF system includes a similar asynchronous mode; however, synchronous operation is likely to provide superior performance under conditions of moderate to high network load.

When operating in synchronous mode, all scanning receivers in a 3G-ALE network change frequency at the same time (to within a relatively small timing uncertainty). It is not necessary that all stations monitor the same calling channel at the same time, however. By assigning groups of network members to monitor different channels in each scanning dwell, calls directed to network member stations will be distributed in time and/or frequency, which greatly reduces the probability of collisions among 3G-ALE calls. This is especially important under conditions of high traffic load. The set of stations that monitor the same channels at the same time is called a *dwell group*.

4.7 Addressing

One of the functions of the subnetwork layer is translation of upper-layer addresses (*e.g.*, IP addresses) into whatever peculiar addressing scheme the local subnet uses. The addresses used in 3G-ALE PDUs are 11-bit binary numbers. In a network operating in synchronous mode, these addresses are partitioned into a 5-bit dwell group number and a 6-bit member number within that dwell group. Up to 32 dwell groups of up to 60 members each are supported (1920 stations per net). Four additional reserved addresses in each group (1111xx) are available for use by stations calling into the network (see Figure 4.3).

When it is desired to be able to reach all network members with a single call, and traffic on the network is expected to be light, up to 60 network member stations may be assigned to the same dwell group. However, this arrangement is subject to calling channel congestion. To support heavier call volume than the single group scheme will support, the network members should be distributed into multiple dwell groups.

4.8 Synchronous dwell structure

The nominal duration of each synchronous dwell is 4 seconds. The timing structure within each synchronous dwell time is as follows.

Slot 0: Tune and Listen Time. During slot 0, radios are switched to the new receiving frequency, couplers are tuned as necessary, and so on. (A calling station will instead tune to the frequency on which it will handshake during that dwell.) Following tuning, every receiver samples a traffic frequency in the vicinity of the new calling channel, attempting to detect traffic. This provides recent traffic channel status before stations get involved in a handshake.

Calling Slots. The remainder of the dwell time is divided into 4 equal-length calling slots. These slots are used for the synchronous exchange of PDUs on calling channels. 800 ms per slot allows for a 26-bit PDU, 70 ms of propagation, and synchronization uncertainty of ± 50 ms.

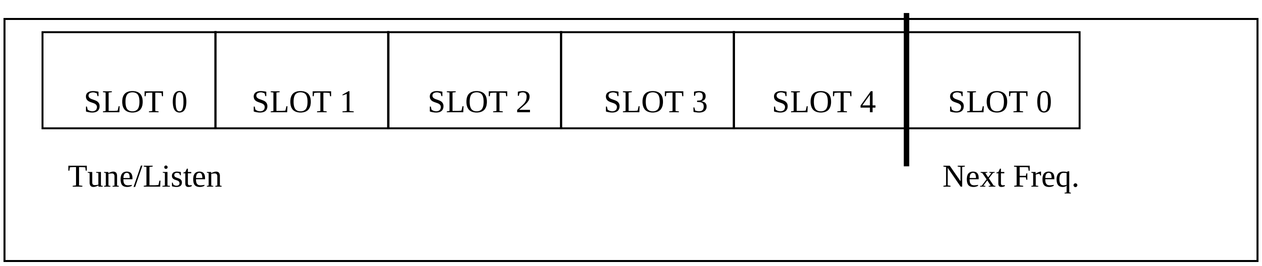

FIGURE 4.3

Synchronous dwell structure

4.9 Synchronous calling overview

The 3G-ALE synchronous calling protocol seeks to find suitable channel(s) for traffic and transition to them as quickly as possible. This minimizes occupancy of the calling channels, which is important in any CSMA system. 3G-ALE calls indicate the type of traffic to be carried (in general terms), and the first traffic channel(s) that will support this grade of service will be used. Note that the default system does not spend time seeking the *best* channels for traffic.

Figure 4.4 and Figure 4.5 show two methods of connection set up and traffic exchange: separate control and traffic frequencies, and identical control and traffic frequencies. Each figure shows connection set-up, a transition to the traffic frequency, traffic set-up, and finally traffic exchange. When directed to establish a link to a prospective *responding* station, the *calling* station will compute the frequency to be scanned by the responding station during the next dwell and randomly (though not uniformly) select a calling slot within that dwell time. During slot 0 of that dwell, the calling station listens to a nearby traffic channel that has recently been free of traffic. (A station with multiple receivers listens to multiple traffic channels during slot 0). If not calling in slot 1, the calling station listens on the calling channel for other calls during the slots that precede its call. If it detects a handshake, it will defer its call until after that handshake. If no other handshake is detected before its chosen slot, the calling station sends a probe PDU (described later) in that slot and listens for a response in the next slot.

When a station receives a probe PDU addressed to it, it responds in the next slot with a handshake PDU. The handshake PDU may indicate a good traffic channel for transmissions to that responding station. If it does, the stations will immediately proceed to the negotiated traffic channel for a traffic set-up handshake, followed by the traffic described in the call.

If the handshake on a channel does not conclude in a transition to traffic channels, the handshake will proceed to the next calling channel in the responding station's scan list during the next dwell. The calling station will again select a slot and start the handshake in this new dwell by sending a probe PDU.

Listen-before-transmit. Every calling station that will send a PDU during a dwell must listen on its intended calling channel during the slot that precedes its transmission (except Slot 1). If it detects a handshake during that slot, it should defer its call. Thus, calls in early slots in a dwell may pre-empt calls in later slots.

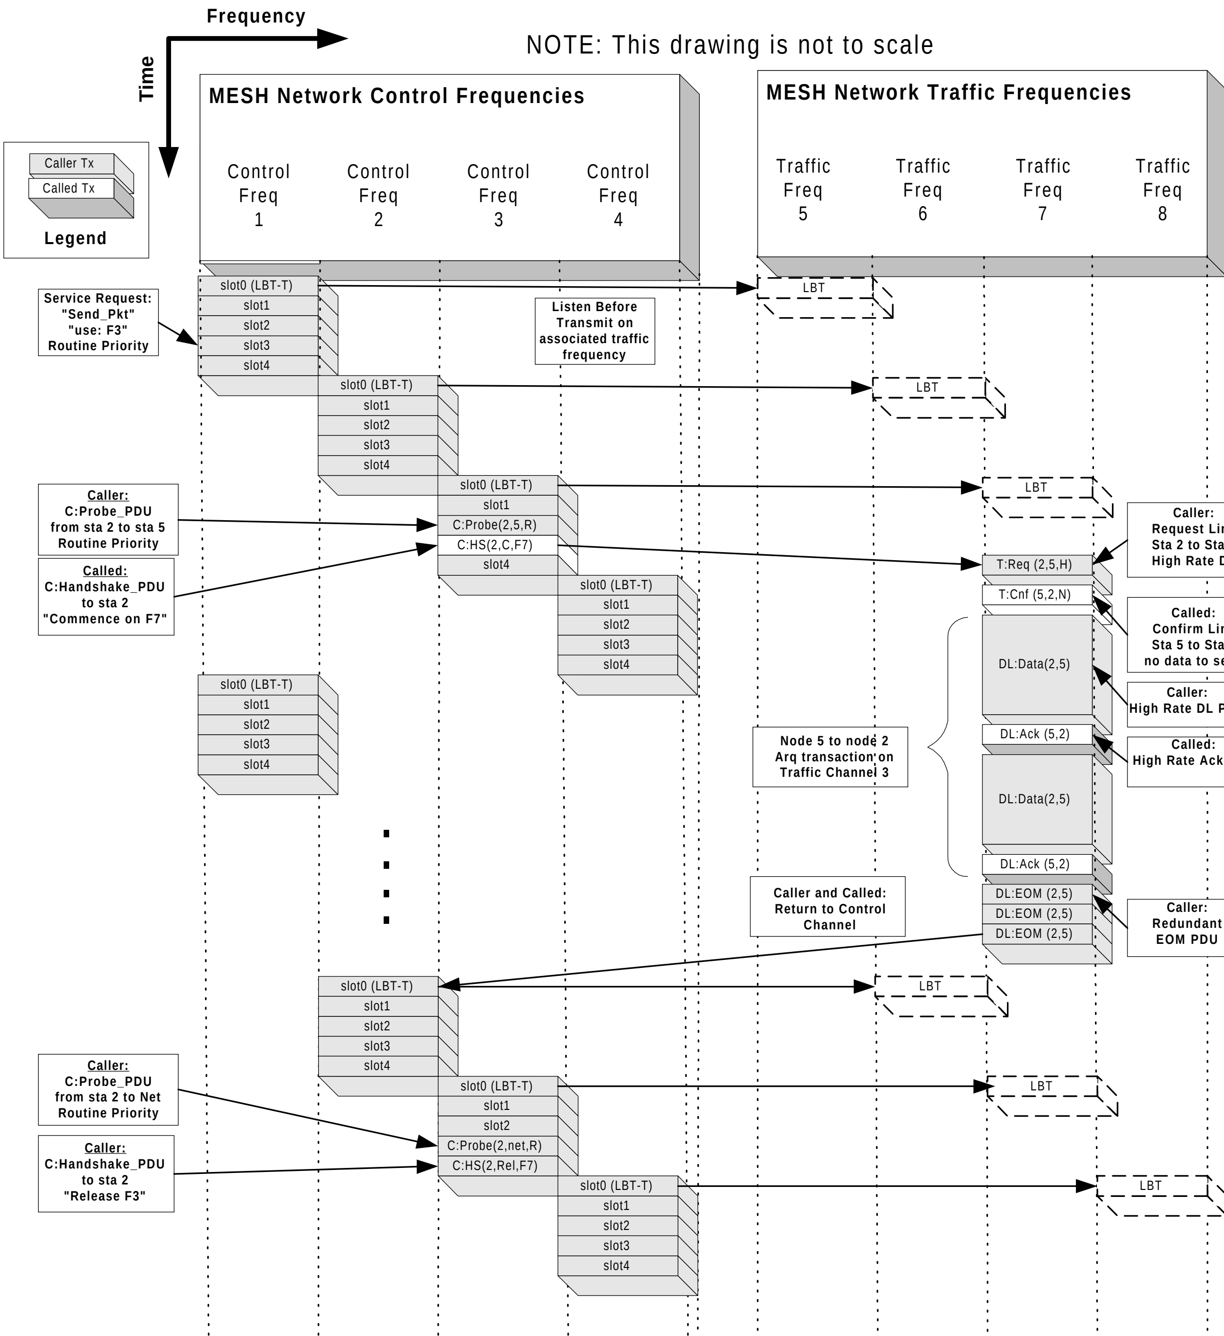

FIGURE 4.4

Concept overview: separate control and data frequencies

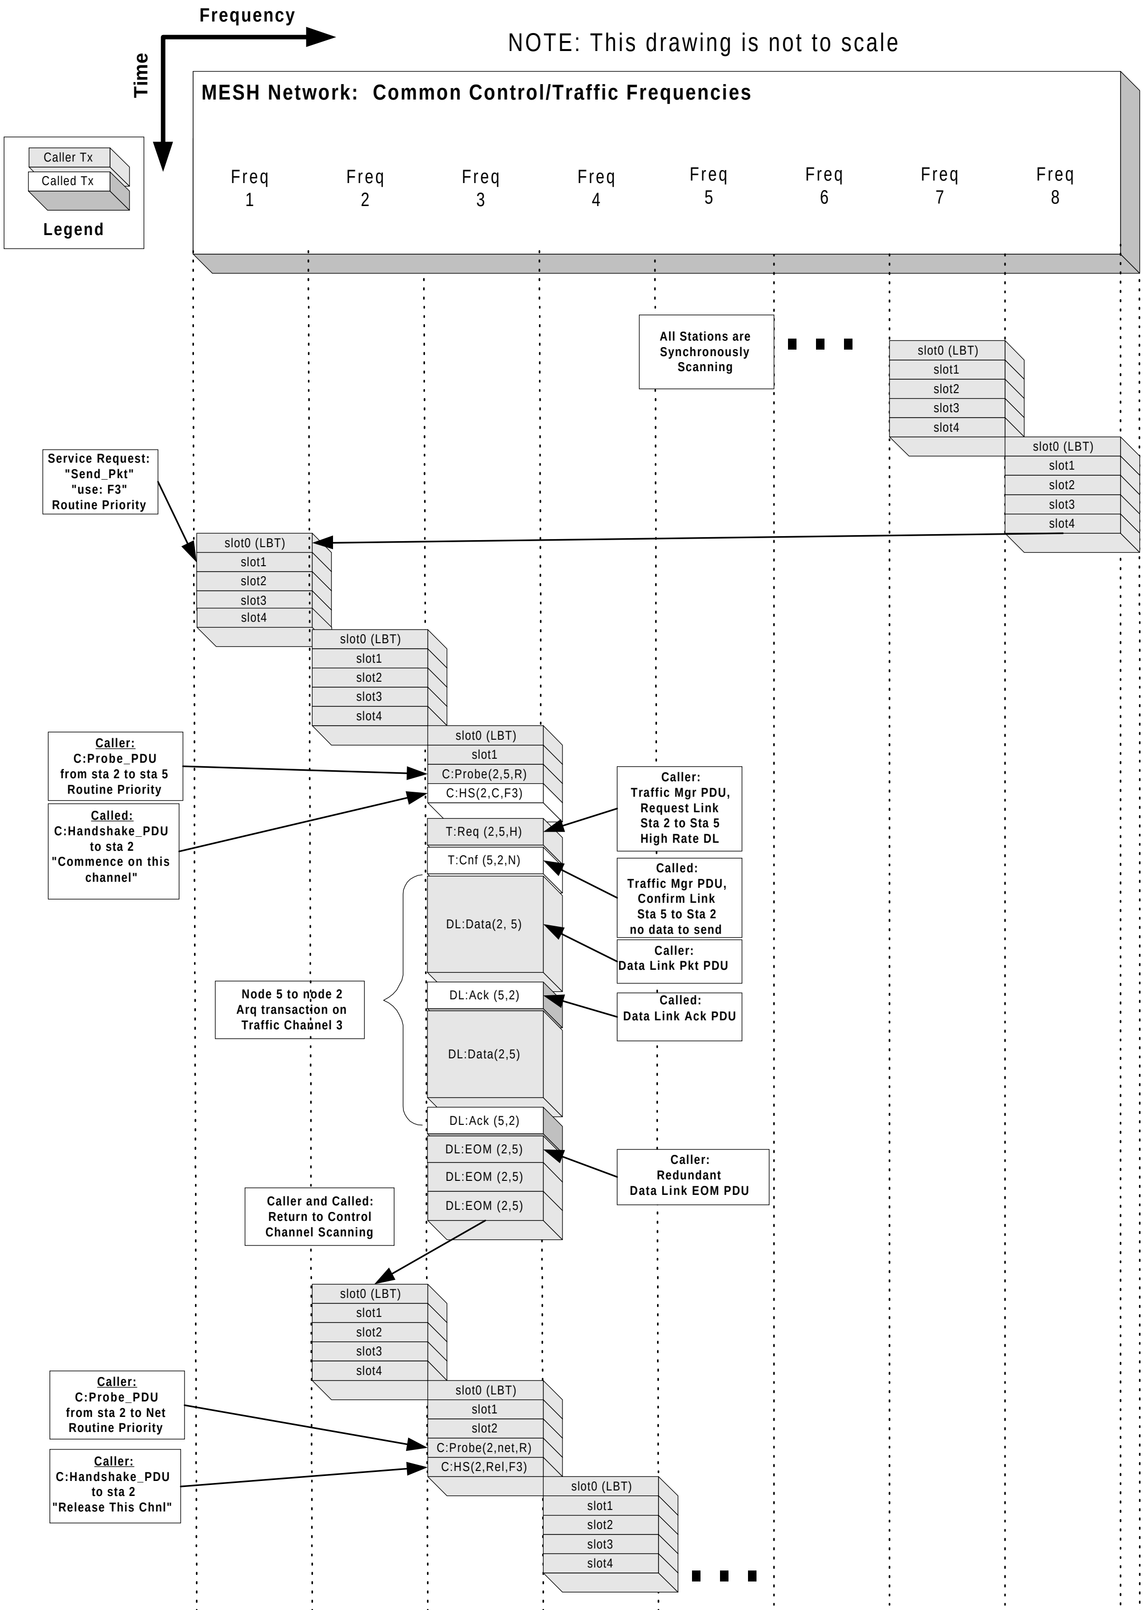

FIGURE 4.5

Concept overview: shared control and data frequencies

Prioritized Slot Selection. The probability of selecting a slot is randomized over all usable slots, but the probabilities for higher-priority calls are skewed toward the early slots while low-priority calls are skewed toward the later slots. Such a scheme will operate reasonably well in all situations, while hard partitioning of early slots for high and late slots for low priorities would exhibit inordinate congestion in crisis and/or routine times. A suggested set of probabilities is shown below for a four-priority implementation:

TABLE 4.2

Probability of slot selection

Call Priority	*Slot 1*	*Slot 2*	*Slot 3*	*Slot 4*
Flash	***50%***	***30%***	***15%***	***5%***
Immediate	***30%***	***50%***	***15%***	***5%***
Priority	***5%***	***15%***	***50%***	***30%***
Routine	***5%***	***15%***	***30%***	***50%***

4.10 Third generation ALE PDUs

The contents of the various PDUs used by the connection manager are shown in Figure 4.6. The PDUs used in one-to-one calling are the probe and handshake PDUs, as noted above. These two key PDUs are discussed below. Readers are directed to the formal draft specification for a description of the more specialized PDUs.

4.10.1 Probe PDU

The probe PDU needs to convey sufficient information to the called station so that it will know whether it wants to respond, and what to listen for during the traffic channel check. The probe PDU must therefore report

- the calling station identification
- the priority of the incoming call
- what resources will be needed if the call is accepted
- what traffic channel quality is required.

The call type field in the probe PDU includes a bit that flags high priority calls, an indication of whether analog voice or a modem will be used for traffic, and the duration of the traffic. Duration is categorized as a “short” data message, a “long” data message, or unbounded duration (virtual circuit). The specific meaning of short versus long is left to the discretion of the network designer. The 3G-DLP uses this field to optimize its choice of data transfer waveform and protocol. The full called station address is not needed in the probe PDU, because the called station group number is implicit in the choice of channel that carries the probe.

Third Generation ALE PDUs

23 April 1998

Probe PDU

	6	3	6	5	4
1 0	Called Member # (not 1111xx)	Call Type	Caller Member #	Caller Group #	CRC

Handshake PDU

	6	3	7	8
0 0	Link ID	Command	Argument (e.g., ch #)	CRC

Notification PDU

	6	3	6	5	4
1 0	1 1 1 1 1 1	Caller Status	Caller Member #	Caller Group #	CRC

Broadcast PDU

	3	3	3	7	8
0 1	1 1 0	Countdown	Call Type	Channel	CRC

Scanning Call PDU

	5	6	5	8
0 1	1 1 1 1 1	Called Member # (not 1111xx)	Called Group #	CRC

Call Type

	0xx is Routine Priority
0 0 0	Short data message
0 0 1	Long data message
0 1 0	Modem circuit
0 1 1	Analog voice circuit
	1xx is High Priority
1 0 0	Short data message
1 0 1	Long data message
1 1 0	Modem circuit
1 1 1	Analog voice circuit

Handshake PDU Commands

	Command	Argument
0 0 0	Continue Handshake	reason
0 0 1	Traffic Channel	channel
0 1 0		
0 1 1		
1 0 0		
1 0 1		
1 1 0	Data	value
1 1 1	Abort Handshake	reason

Caller Status

0 0 0	Nominal
0 0 1	
0 1 0	
0 1 1	
1 0 0	
1 0 1	
1 1 0	Commencing EMCON
1 1 1	Leaving network

FIGURE 4.6

Third generation ALE PDUs

4.10.2 Handshake PDU

The Handshake PDU is used by both calling and called stations. It is sent only after a Probe PDU has established the identities of both stations in one-to-one link establishment, as well as the key characteristics of the traffic that will use the link. The Link ID field contains a 6-bit hash of the 11-bit addresses of the calling and responding stations. The commands carried in Handshake PDUs include the following:

Continue Handshake	The handshake will continue for at least another two-way handshake and will proceed to the next assigned called station dwell frequency. The argument is a Reason code (*e.g.*, Poor Propagation or Channel Busy).
Traffic Channel	This is the final command sent on a calling channel. The argument is the channel number on which the responding station will listen for traffic. Following this command, both stations proceed to that traffic channel.
Abort Handshake	This command immediately terminates the handshake and needs no response. It is analogous to the TWAS preamble in second-generation ALE.
Data	This command is used only in special-purpose protocols. The argument carries previously requested data.

4.11 One-to-one link establishment

The one-to-one linking protocol identifies a frequency for traffic use relatively quickly (*i.e*., within a few seconds), and minimizes channel occupancy during this link establishment process. It will conclude the link establishment process as soon as a suitable frequency has been identified, and makes no attempt to find the best available frequency.

4.12 Multicast calls

A multicast call employs a reserved member number in each affected dwell group (similar to a 2G-ALE Net Call). 3G-ALE controllers must be programmed to recognize multicast addresses to which they subscribe. That is, multicast addresses are reserved within a network, not in the protocol standard.

The multicast protocol works as follows:

1. A Probe PDU is sent as usual, but it contains a multicast responding-station address, which suppresses responses by the responding station(s). The Call Type field specifies whether the multicast link will carry modem or analog voice traffic when established.

2. The *caller* then sends a Traffic Channel PDU in the immediately following time slot that directs the responding stations to a traffic channel where they are to listen for the type of traffic specified in the call. This completes the multicast calling protocol.

4.13 Other CM PDUs

Other types of connection manager PDUs are provided to support broadcast calls, notification calls, and asynchronous scanning calls. Readers are referred to the formal draft specification for a description of these connection manager services.

4.14 Traffic manager

Traffic set-up is accomplished using the BW1 40-bit burst waveform. The traffic management (TM) protocol is used to co-ordinate traffic exchanges on connections established using the connection management (CM) protocol. Following the end of the connection set-up (CSU) phase in which a connection is established, the stations participating in the connection enter the traffic set-up (TSU) phase in which the traffic management protocol is used to establish a traffic link on which traffic can be delivered.

Once a connection has been established, the stations participating in it have determined:

1. the identities of the stations intended to participate in the connection;
2. the connection topology: point-to-point, multicast, or broadcast;
3. the link mode: packet or circuit (“hard link”)
4. the HF frequency (or “traffic channel”) that will be used for signaling within the connection.

In addition, the initiating station knows it can transmit a traffic management PDU in the first transmit time-slot of the TSU phase.

During the TSU phase, the participating stations exchange TM PDUs in order to determine:

1. which data link protocol(s), waveform(s), and/or baseband modulation formats will be used to deliver traffic on the connection;
2. the priority of the traffic to be delivered;
3. the fine time synchronization required for the MIL-STD-188-141B Appendix C High-Rate and Low-Rate Data Link Protocols, on traffic links established for packet traffic.

TABLE 4.3
Traffic management (TM) PDU format

field name	size (bits)	values	description
Protocol	2	3 (fixed)	distinguishes TM PDUs from HDL_ACK and HDL_EOM PDUs (equal to 11 binary)
Priority	2	In all TM_REQ and some TM_CONF PDUs (i.e., TM PDUs having Type = TM_REQUEST or TM_CONFIRM), indicates the priority level of the traffic (if any) that the sender of the PDU intends to send on the traffic link once it is established.	
		0	Flash: highest priority
		1	Immediate
		2	Priority
		3	Routine: lowest priority
Dest Addr	11	any	address of the station to which this PDU is being sent. In the PDUs exchanged to manage a broadcast or multicast traffic link, this address may be the all-ones broadcast address (“BCaddr”) or a net multicast address (“MCaddr”) with all ones in the dwell group number subfield.
Source Addr	11	any	address of the station that is sending this PDU. Is always the station address of a single station — never a multicast or broadcast address.
Type	3	Type of PDU, indicating its role in the Traffic Management protocol. Note that the state diagrams and other materials refer to, for instance, a “TM_REQ PDU”; this is a Traffic Management PDU whose Type field value is 0 (TM_REQUEST).	
		0	TM_REQUEST: A PDU with Type = TM_REQUEST is sent in order to request that a traffic link be established between the station sending the TM_REQUEST and the other stations specified by the PDU’s destination address.
		1	TM_CONFIRM: A PDU with Type = TM_CONFIRM is sent in response to a received TM_REQUEST PDU, to confirm the sender’s readiness to participate in a traffic link.
		2	TM_TERM: a PDU with Type = TM_TERM is sent in order to terminate the station’s participation in a traffic link (during or after link establishment), and when sent by the link master, to terminate the link as a whole.
		3..7	reserved
Argument	3	variant field whose usage and meaning depend on the value of the Type field.	
CRC	8	any	8-bit Cyclic Redundancy Check (CRC) computed across the remaining 32 bits of each Traffic Management PDU, using the generator polynomial $X^8 + X^7 + X^4 + X^3 + X^1 + 1$.

If the traffic link is to be used for multicast communication, the participating stations conduct a roll-call procedure to determine which of the stations in the multicast group received the CSU signaling and are now present on the traffic frequency.

When traffic exchanges have been completed on a traffic link, the traffic management protocol is used to co-ordinate the participating stations’ departure from the traffic link.

Behavioral descriptions of the traffic management protocol refer to three kinds of PDUs: TM_REQ, TM_CONF, and TM_TERM. These PDUs all have the format shown in TABLE 4.3, and are distinguished from one another by the values of their Type fields:

- a “TM_REQ PDU” is a traffic management PDU having Type = TM_REQUEST (0)
- a “TM_CONF PDU” is a traffic management PDU having Type = TM_CONFIRM (1)
- a “TM_TERM PDU” is a traffic management PDU having Type = TM_TERM (2).

Figure 4.7 presents a state diagram defining the behavior of the traffic management protocol in setting up traffic links for packet traffic, and illustrates the state diagram notation used in the protocol specification. In the state diagram, each state transition is labeled with an event, an optional condition, and zero or more actions. This indicates that the state transition occurs whenever the event occurs *and* the condition obtains (is TRUE), causing the associated actions to be performed. In the diagram,

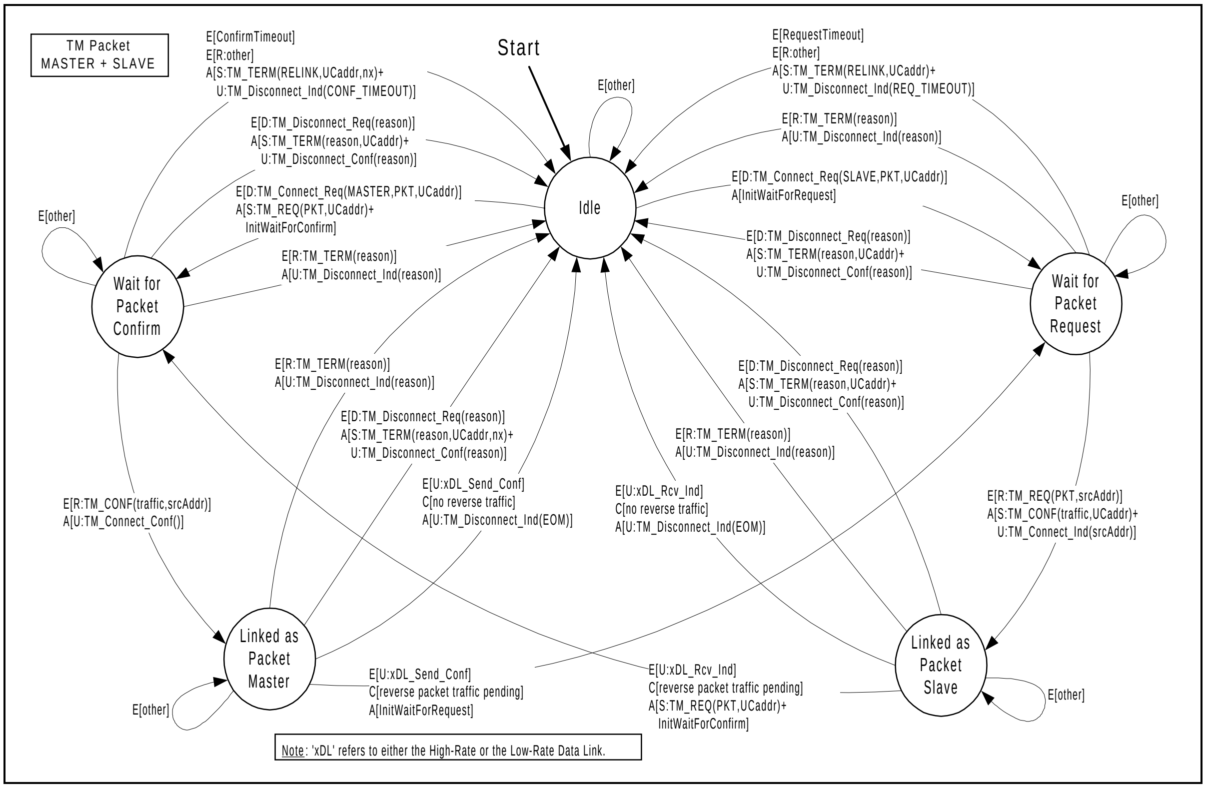

FIGURE 4.7

Traffic manager state diagram (excerpt)

- the name of each event is shown in brackets preceded by the letter ‘E’;
- the description of each condition is shown in brackets preceded by the letter ‘C’; and
- the names of the actions associated with a transition are shown in brackets preceded by the letter ‘A’.

Where a transition is labeled with two or more events, this indicates that the transition occurs whenever any of the events occurs. The prefix ‘R:’ in the name of an event indicates that the event is the receipt of a PDU from the remote station. ‘D:’ indicates that the event is an Traffic Management service primitive passed *down* to TM from a higher-layer entity; ‘U:’ indicates a lower-layer service primitive passed *up* to TM from a lower-layer entity. Similar conventions are used in naming actions, with the addition that the prefix ‘S:’ in the name of an action indicates that the action sends a PDU to the remote station.

4.15 Data link protocols

Two data link protocols are provided: one for large messages and/or good channel conditions, and a second for short messages and/or poor channel conditions. Both data link protocols use memory combining and do not require data rate adaptation. This greatly simplifies the protocols while dramatically increasing throughput under nearly all channel conditions and signal to noise ratios.

4.15.1 High-rate data link protocol

The high-rate data link protocol (HDL) is used to provide acknowledged point-to-point delivery of datagrams from a transmitting station to a receiving station across an already established HF link, with selective retransmission (ARQ) of data received in error. The datagram passed to the high-rate data link for delivery is an ordered sequence of up to 7,634,944 8-bit data bytes (octets). The high-rate data link protocol is best suited to delivering relatively large datagrams under good to fair HF channel conditions. By contrast, the low-rate data link protocol described below provides better performance for all datagram lengths under fair to very poor HF channel conditions, and under all channel conditions for short datagrams.

Data transfer by HDL begins after the stations have already established the data link connection in the traffic set-up phase, in so doing negotiating the fact that HDL will be used (as opposed to LDL or some other mechanism), the number of data packets to be sent in each HDL_DATA PDU, and the precise time synchronization of data link transmissions.

In an HDL data transfer, the sending station and the receiving station alternate transmissions in the manner depicted in Figure 4.8, the sending station transmitting HDL_DATA PDUs containing payload data packets, and the receiving station transmitting HDL_ACK PDUs containing acknowledgements of the data packets received without errors in the preceding HDL_DATA PDU. If either station fails to receive a PDU at the expected time, it sends its own next outgoing PDU at the same time as if the incoming PDU had been received successfully. The times at which the burst waveforms conveying HDL_DATA, HDL_ACK, and HDL_EOM PDUs may be transmitted are determined precisely by the initial data link timing established during the traffic set-up phase.

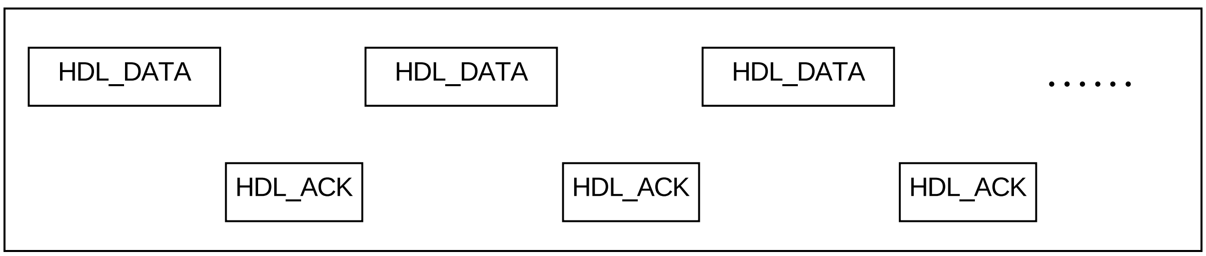

FIGURE 4.8.

HDL data transfer overview

The end of a data transfer is reached when the sending station has transmitted HDL_DATA PDUs containing all of the payload data in the delivered datagram, and the receiving station has received these data without errors and has acknowledged their successful delivery. When the sending station receives an HDL_ACK PDU indicating that the entire contents of the datagram have been delivered successfully, it sends an HDL_EOM PDU repeated as many times as possible within the duration of an HDL_DATA PDU, starting at the time at which it would have otherwise transmitted the next HDL_DATA PDU, to indicate to the receiving station that the data transfer will be terminated. This link termination scenario is depicted in Figure 4.9.

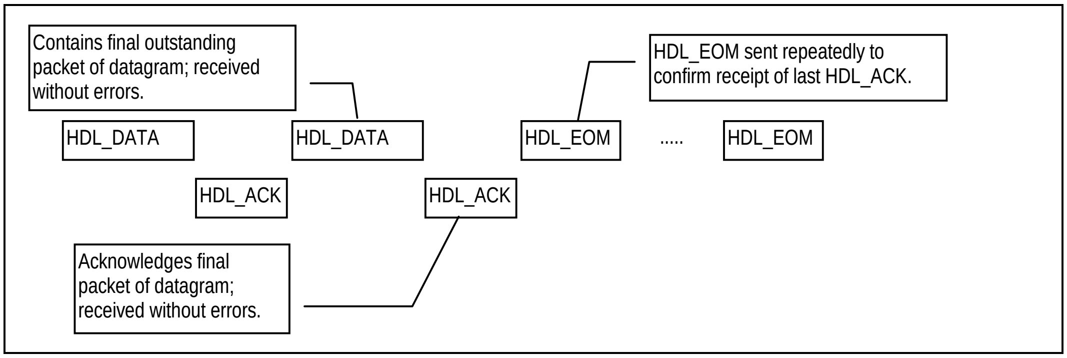

FIGURE 4.9

HDL link termination scenario overview

4.15.2 High-rate data link PDUs

Figure 4.10 depicts the format and contents of the high-rate data link's PDUs. Each HDL_DATA PDU is a sequence of 24, 12, 6, or 3 data packets, in which each packet is composed of 1881 bits of payload data (1864 bits of user data plus a 17-bit sequence number added by the data link). During the traffic set-up phase, the user process determines the number of data packets per HDL_DATA PDU so as to deliver the user data efficiently, shortening the HDL_DATA PDU whenever the entire datagram is short enough to fit within the shortened PDU. Once it is determined, the number of data packets per HDL_DATA PDU for the current datagram delivery is communicated to the receiving station in the traffic set-up sequence (see Section 4). Thereafter, every HDL_DATA PDU contains the same number of data packets until the entire datagram has been delivered. The BW2 waveform is used to transmit each HDL_DATA PDU; further description of BW2 processing is provided below.

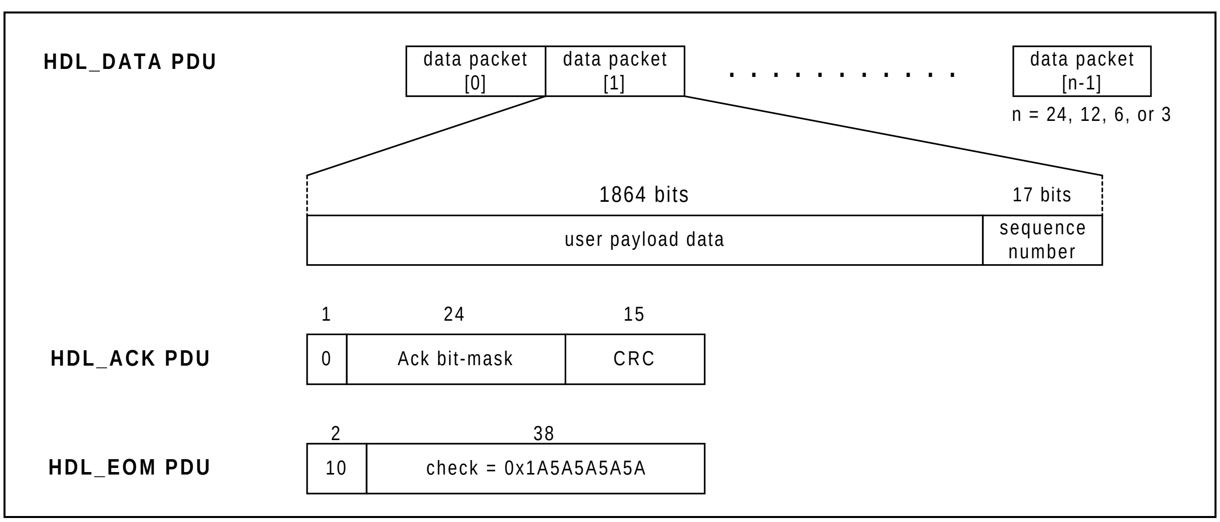

FIGURE 4.10

High-rate data link PDUs

The HDL_ACK PDU is used to convey data acknowledgements from the receiving station to the sending station. Each HDL_ACK PDU contains acknowledgements for the immediately preceding HDL_DATA PDU sent in the opposite direction; each bit in the 'Ack bit-mask' field acknowledges a single corresponding data packet from the HDL_DATA PDU. The HDL_ACK PDU contents are protected by a 15-bit CRC.

The HDL_EOM PDU is transmitted in the forward direction, in place of an HDL_DATA PDU, when the sending station receives an error-free HDL_ACK PDU indicating that the entire user datagram has been delivered to the receiving station without errors.

BW1 is used to transmit both the HDL_ACK and HDL_EOM PDUs. The marker bits at the beginning of each PDU are used to distinguish the two kinds of PDUs.

4.15.3 High-rate ARQ processing

Figure 4.11 depicts the manner in which an HDL_DATA PDU is incorporated into a BW2 burst transmission. Each data packet in the HDL_DATA PDU is extended by appending to it a 32-bit CRC, followed by an encoder flush sequence consisting of seven zero bits. The resulting sequence of extended packets is FEC-encoded using a ¼-rate convolutional encoder. The encoder produces four output bits, $Bitout_0$ .. $Bitout_3$, for each input bit. As each packet is encoded, the bits from each encoder output are accumulated into a block, resulting in four blocks of coded bits, $EBlk_0$ .. $EBlk_3$. Each time a data packet is transmitted in an HDL_DATA PDU, only one of the four blocks of coded bits from the packet is transmitted, starting with $EBlk_0$ the first time. (The original contents of the packet can be recovered from any single block of coded bits that is received without errors.) Each time the data packet cannot be decoded without errors and must be retransmitted, a different block of coded bits is transmitted; blocks are transmitted in the order $EBlk_0$, $EBlk_1$, $EBlk_2$, $EBlk_3$, $EBlk_0$, ... The transmission of different blocks of coded bits for each packet, in successive transmissions of the packet, provides additional information that can be used in decoding the packet.

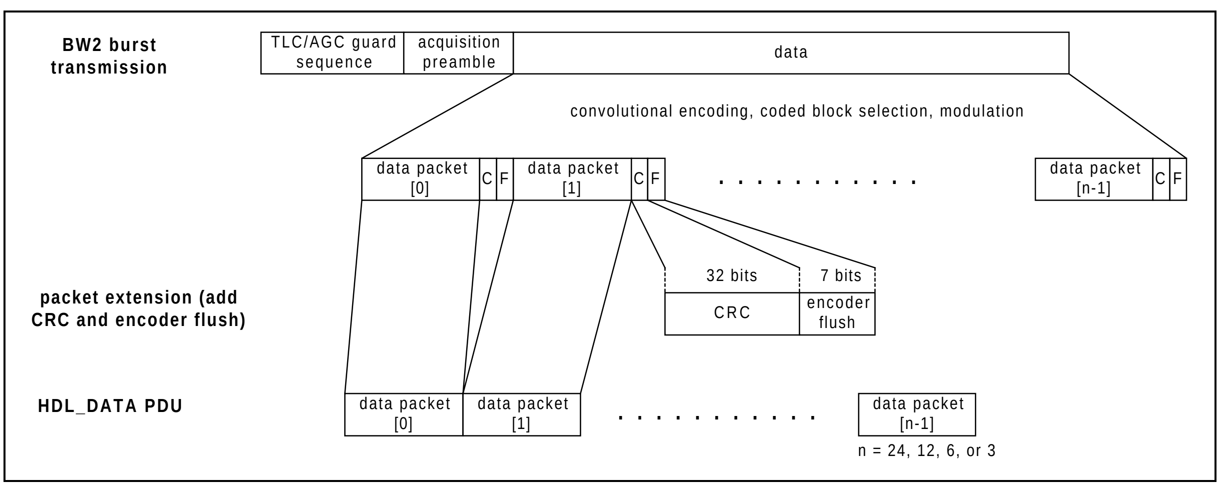

FIGURE 4.11

BW2 encoding and modulation of HDL_DATA PDU

The sequence of coded bits is modulated using a modulation process similar to that of the 110A serial tone waveform at 4800 bits per second. This results in a sequence of unknown/known symbol frames, each consisting of 32 unknown symbols (8-ary PSK symbols carrying three bits each, Gray-coded), followed by 16 known symbols. A TLC/AGC guard sequence and a sequence of 64 known symbols used for initial equalizer training are prepended to the beginning of the frame sequence. Note that no acquisition preamble is required in the BW2 waveform, since the precise time of arrival of each BW2 transmission is known once the traffic set-up handshake is used to establish data link timing.

4.16 Low-rate data link protocol

The low-rate data link protocol (LDL) is used to provide reliable acknowledged point-to-point delivery of datagrams from a transmitting station to a receiving station across an already-established HF link. The datagram passed to the low-rate data link protocol entity for delivery is an ordered sequence of up to 16,384,000 8-bit data bytes (octets). The low-rate data link protocol provides better performance than does the high-rate data link protocol for all datagram lengths under fair to very poor HF channel conditions, and under all channel conditions for short datagrams.

Data transfer by LDL begins after the traffic management sublayer has already established the data link connection, in so doing negotiating the fact that LDL will be used (as opposed to HDL), and the precise time synchronization of data link transmissions. In an LDL data transfer, the sending station and the receiving station alternate transmissions in the manner depicted in Figure 4.12, the sending station transmitting LDL_DATA PDUs containing payload data packets, and the receiving station transmitting LDL_ACK PDUs each containing an acknowledgement of whether or not the data packet in the preceding LDL_DATA PDU was received without error. If either station fails to receive a PDU at the expected time, it sends its own next outgoing PDU at the same time as if the incoming PDU had been received successfully. The times at which the burst waveforms conveying LDL_DATA, LDL_ACK, and LDL_EOM PDUs may be emitted are determined precisely by the initial data link timing established during the traffic set-up phase.

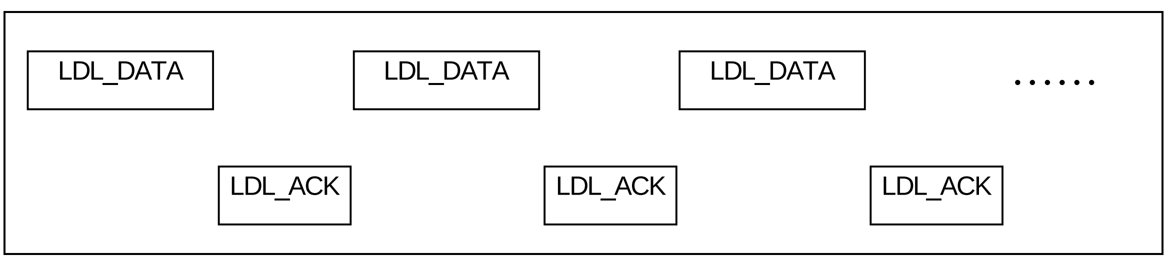

FIGURE 4.12

LDL data transfer overview

The end of a data transfer is reached when the sending station has transmitted LDL_DATA PDUs containing all of the payload data in the delivered datagram, and the receiving station has received these data without errors and has acknowledged their successful delivery. When the sending station receives an LDL_ACK PDU indicating that the entire contents of the datagram have been delivered successfully, it sends an LDL_EOM PDU repeated as many times as possible within the duration of an LDL_DATA PDU, starting at the time at which it would have otherwise transmitted the next LDL_DATA PDU, to indicate to the receiving station that the data transfer will be terminated. This link termination scenario is depicted in Figure 4.13.

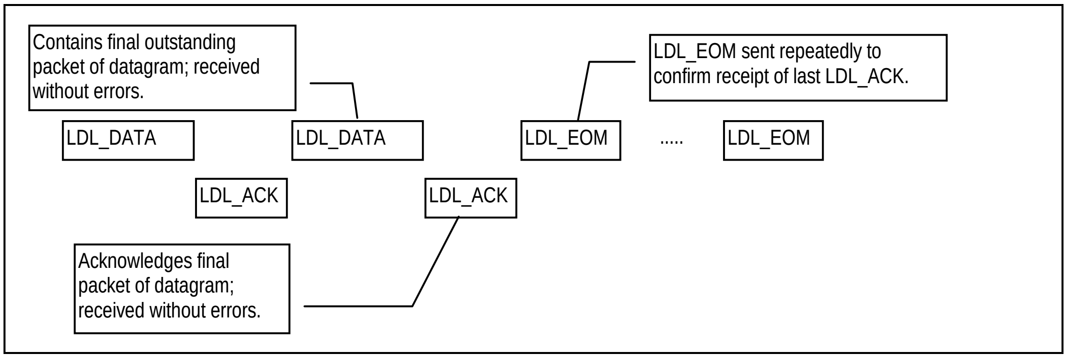

FIGURE 4.13

LDL link termination scenario overview

4.16.1 Low-rate data link PDUs

Figure 4.14 depicts the format and contents of the low-rate data link PDUs. Each LDL_DATA PDU carries a single data packet composed of payload data (512, 256, 128, or 64 *bytes* (octets) of user data) followed by a 17-bit sequence number and an 8-bit control field (presently unused) added by the low-rate data link. During the traffic set-up phase, the user process determines the number of data bytes per LDL_DATA PDU so as to deliver the user data efficiently, shortening the LDL_DATA PDU whenever the entire datagram is short enough to fit within the shortened PDU. Once it is determined, the number of data bytes per LDL_DATA PDU for the current datagram delivery is communicated to the receiving station in the traffic set-up sequence. Thereafter, every transmitted LDL_DATA PDU contains the same number of data bytes until the entire datagram has been delivered.

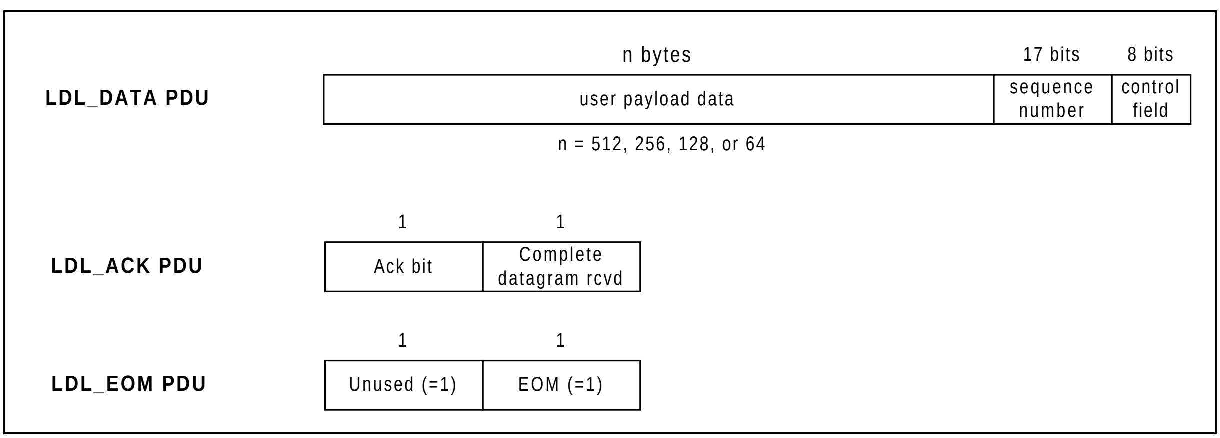

FIGURE 4.14

Low-rate data link PDUs

The LDL_ACK PDU is used to convey data acknowledgements from the receiving station to the sending station. Each LDL_ACK PDU contains acknowledgements for the immediately preceding LDL_DATA PDU sent in the opposite direction; the single bit in the 'Ack bit' field acknowledges the single data packet in the LDL_DATA PDU. The 'Complete datagram rcvd' bit is set when the receiving station determines that it has received all of the contents of the datagram without errors, so that the data link transfer can be ended. The LDL_ACK PDU is transmitted using the very robust BW4 waveform. Due to the robustness of the waveform, no CRC is included in the PDU.

The LDL_EOM PDU is transmitted in the forward direction, in place of an LDL_DATA PDU, when the sending station receives an error-free LDL_ACK PDU indicating that the entire user datagram has been delivered to the receiving station without errors. This PDU is also transmitted using the BW4 waveform. LDL_EOM PDUs are distinguished from LDL_ACK PDUs by context: any BW4 transmission in the *forward* direction of a low-rate data link transfer is an LDL_EOM PDU.

4.16.2 Low-rate ARQ processing

Figure 4.15 depicts the manner in which an LDL_DATA PDU is incorporated into a BW3 burst transmission. The LDL_DATA PDU is extended by appending to it a 32-bit CRC, followed by an encoder flush sequence consisting of seven zero bits. The resulting data sequence is FEC-encoded using a ½-rate convolutional encoder. The encoder produces two output bits, $Bitout_0$ and $Bitout_1$, for each input bit. As each packet is encoded, the bits from each encoder output are accumulated into a block, resulting in two blocks of coded bits, $EBlk_0$ and $EBlk_1$. Each time a data packet is transmitted in an LDL_DATA PDU, only one of the two blocks of coded bits from the packet is transmitted, starting with $EBlk_0$ the first time. (The original contents of the packet can be recovered from any single block of coded bits that is received without errors.) Each time the data packet cannot be decoded without errors and must be retransmitted, a different block of coded bits is transmitted; blocks are transmitted in the order $EBlk_0$, $EBlk_1$, $EBlk_0$, $EBlk_1$, ... The transmission of different blocks of coded bits for each packet, in successive transmissions of the packet, provides additional information that can be used in decoding the packet.

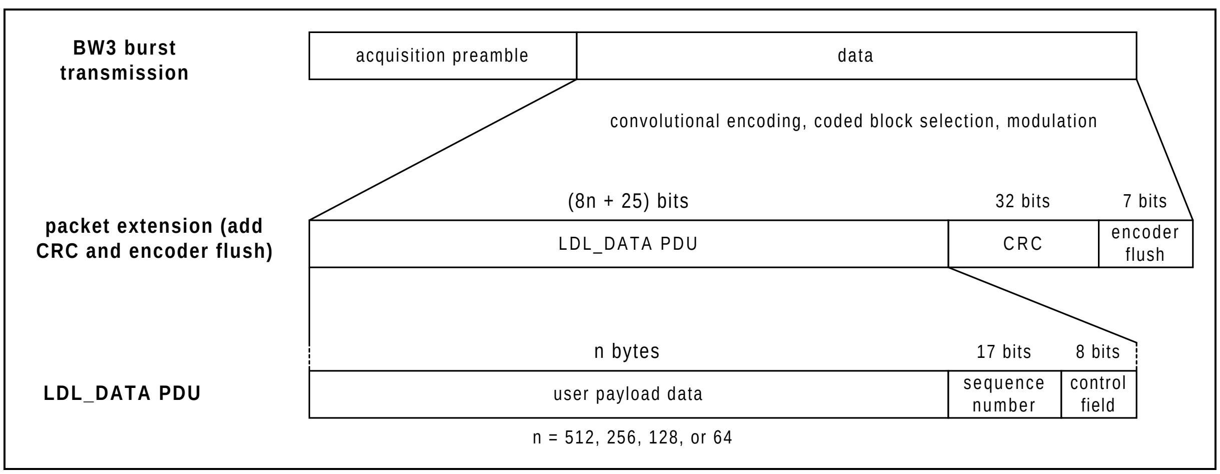

FIGURE 4.15

BW3 encoding and modulation of LDL_DATA PDU

The sequence of coded bits is interleaved using a convolutional block interleaver similar to that of MIL-STD-188-110A. The interleaved bit sequence is then modulated using a modulation process similar to that of the 110A serial tone waveform at 75 bits per second. This results in a sequence of 16-ary orthogonal Walsh frames, each composed of 16 PSK symbols, with each 16-ary Walsh frame representing the values of four coded bits fetched from the interleaver.

An acquisition preamble of 640 8-ary PSK symbols is prepended to the beginning of the Walsh frame sequence. This sequence is used for initial channel estimation but need not be used for synchronization, since the precise time of arrival of each BW3 transmission is known once the traffic set-up handshake is used to establish data link timing.

4.6.3 Data link performance

Data link protocol performance is typically defined and measured in terms of the average throughput in bits per second. The throughput achieved is dependent on many factors, including HF channel conditions, both short term and long term, as well as datagram size. Protocol parameters, which may be selected by the user or automatically adapted, can also play a large role in determining throughput. These parameters may include modem transmission rates, frame or packet size, link turn around times, *etc*.

Data link protocol performance for the high-rate data link protocol (HDL) and the low-rate data link protocol (LDL) are presented in this section. Also presented in this section, for comparison, is the measured performance of the U.S. Federal Standard 1052 data link protocol (FS-1052) [2]. FS-1052 uses the U.S. MIL-STD-188-110A serial tone modem waveform. The autobaud capability of the modem is used extensively as the protocol adapts the baud rate and interleaver settings to the HF channel conditions. Reference [3] presents an overview of the protocol as well as some performance data. Reference [4] presents the asymptotic performance of the data link protocol.

The throughput rates presented in the following figures account for the entire time spent on the Traffic frequency after the completion of a successful call set-up (CSU), including the time for the traffic set-up handshake and the time for the transfer of the message. For HDL and LDL, the traffic set up times are based on the BW1 handshake. In the case of FS-1052, the data link's own handshake timing is included instead of the BW1 handshake timing.

Figures 4-16 through 4-21 present simulated throughput performance for HDL and LDL, and compare their performance to measured 1052 performance for 50, 500, and 50K byte files. Each data link protocol's performance is given for the additive gaussian noise (AWGN) and CCIR Poor HF channel conditions. AWGN refers to a single non-fading path with additive white Gaussian noise. CCIR Poor refers to dual fading paths, separated by 2 ms, each path with a 2 sigma fading bandwidth of 1 Hz.

In comparing the HDL and LDL throughput performance curves against FS-1052, one must be careful to understand some of the basic differences between the two data link protocol techniques. The throughput curves for FS-1052 do not contain a two-way handshake between the two stations in transfers of smaller files (*i.e.,* 50 and 500 bytes); instead, FS-1052 sends a single one-way herald at the beginning of the data transfer. The handshake in the front of the data transfer can dominate throughputs for small files. For larger size files, FS-1052 uses a two-way Call-Response handshake and link termination as part of the file transfer, and therefore allows for a better comparison between techniques. Also note that FS-1052's user-selected forward bit-rate setting can bias its indicated throughput for smaller files. A high initial bit rate setting helps for high SNR at the cost of reduced throughput for low SNR conditions. The FS-1052 data presented here uses 1200 b/s as the initial forward bit rate setting. Finally, overall performance of either system is influenced by the call set-up and traffic set-up mechanisms as well as by the data link protocol. It is hard to compare the performance of the data link protocols in isolation because the data link protocol is not always the limiting factor in the performance of the system as a whole.

In comparing HDL to LDL some interesting observations can be made. HDL has been optimized for higher throughputs for fair to good channel conditions. LDL has been optimized for better operation under severe to fair channel conditions through its choice of its underlying waveform. LDL's orthogonal signaling allows for better throughput performance at lower signal to noise ratios than does HDL's 8-ary PSK signaling. Also, LDL performs better for small message sizes because it incurs less overhead than HDL at these sizes. HDL incurs less protocol overhead for larger files because of its better ratio of forward- to back- channel transmission times. Therefore, HDL is more efficient at the high end of the curves for larger file sizes.

It is important to see that in all of the conditions presented, either HDL or LDL provides throughput performance at least roughly equal to that of FS-1052; in many conditions, the performance of HDL or LDL is dramatically superior. In many conditions, HDL or LDL achieves throughput performance equal to that of FS-1052 at much lower SNR. This fact allows the delivery of equivalent FS-1052 throughput performance at a substantial reduction in radio transmit power. Additionally, for good SNR conditions, HDL can deliver substantially higher throughputs than can FS-1052.

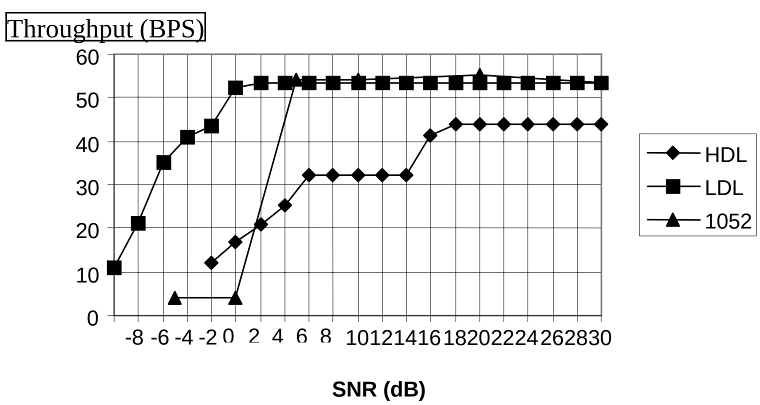

FIGURE 4.16

Gaussian, 50-byte message

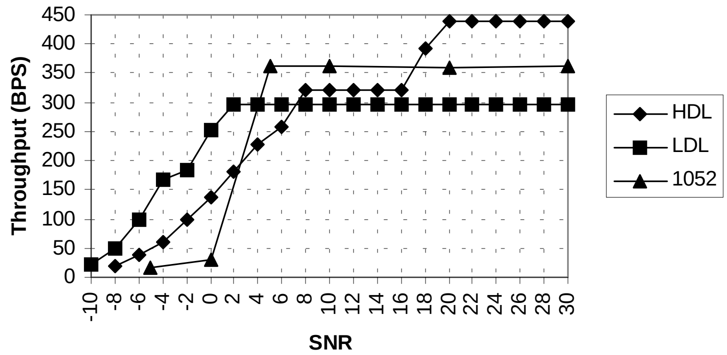

FIGURE 4.17

Gaussian, 500-byte message

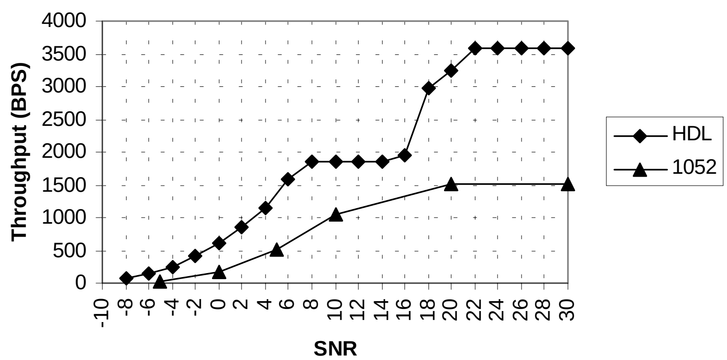

FIGURE 4.18

Gaussian, 50-kbyte message

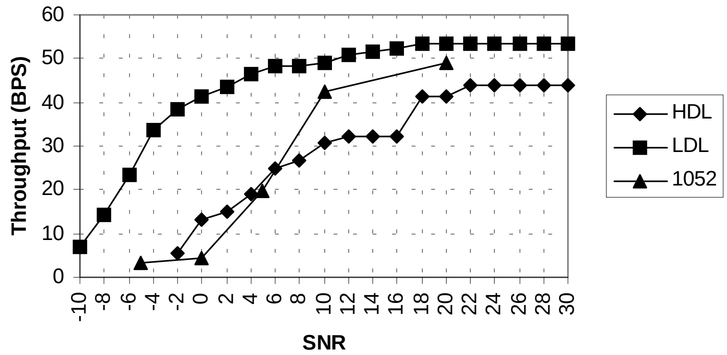

FIGURE 4.19

CCIR Poor, 50-byte message

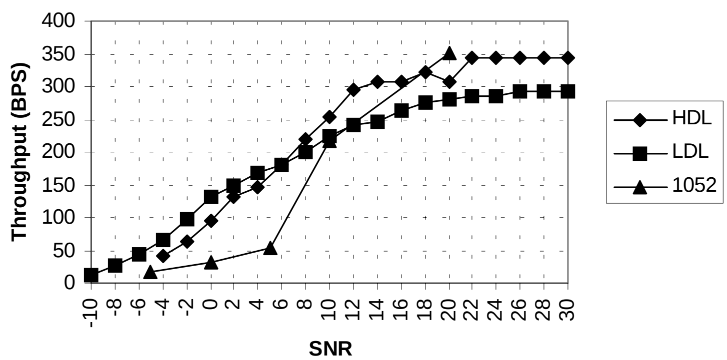

FIGURE 4.20

CCIR Poor, 500-byte message

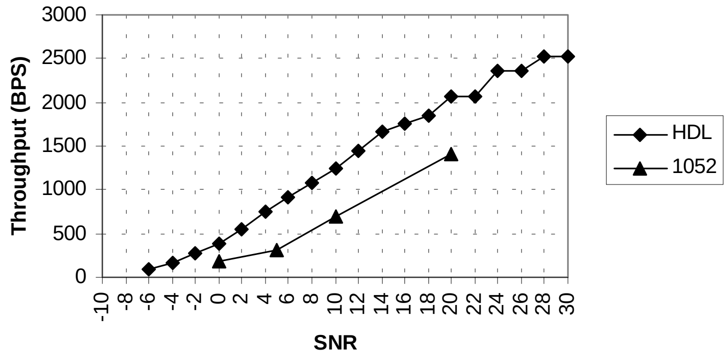

FIGURE 4.21

CCIR Poor, 50-kbyte message

4.17 Circuit link protocol

4.17.1 Overview

The circuit link controller monitors and coordinates traffic on an established circuit link. It provides a simple listen-before-transmit access control mechanism:

- Transmission of new outgoing traffic is inhibited whenever the CLC detects that the circuit link is busy, due to either traffic being received from another station, or traffic currently being transmitted by the local station.
- At the end of each outgoing or incoming traffic transmission, the CLC continues to inhibit transmission of new outgoing traffic for the duration of a backoff interval.

In addition, the CLC provides a traffic timeout indication when an interval of a specified duration elapses in which no outgoing or incoming traffic is detected on the circuit link, allowing the traffic link to be terminated when no longer required.

The CLC is employed only on simplex circuit links (in which all signaling by participating stations is transmitted on a single frequency).

4.17.2 Behavior

The state diagram in Figure 4.22 specifies the behavior of the circuit link controller. In its Idle state, the CLC does not monitor link traffic or control access to the link. When a circuit link is established, the CLC is placed in its Ready state and begins to monitor traffic on the circuit link. From Ready it proceeds to its Transmit or its Receive state, respectively, when outgoing or incoming traffic is detected. When the traffic ends, the CLC proceeds into its Backoff state where it waits for the duration of a backoff interval before returning to its ready state. If incoming signal presence is lost during reception of incoming modem signaling, the CLC enters its “Signal Reacq” state, where it remains until either incoming signal presence is reacquired, or a “ReacqTimeout” event occurs causing the CLC to decide that the incoming traffic has ended and proceed to its backoff state. Note that on the occurrence of any event not shown here, the CLC will take no action and remain in its current state.

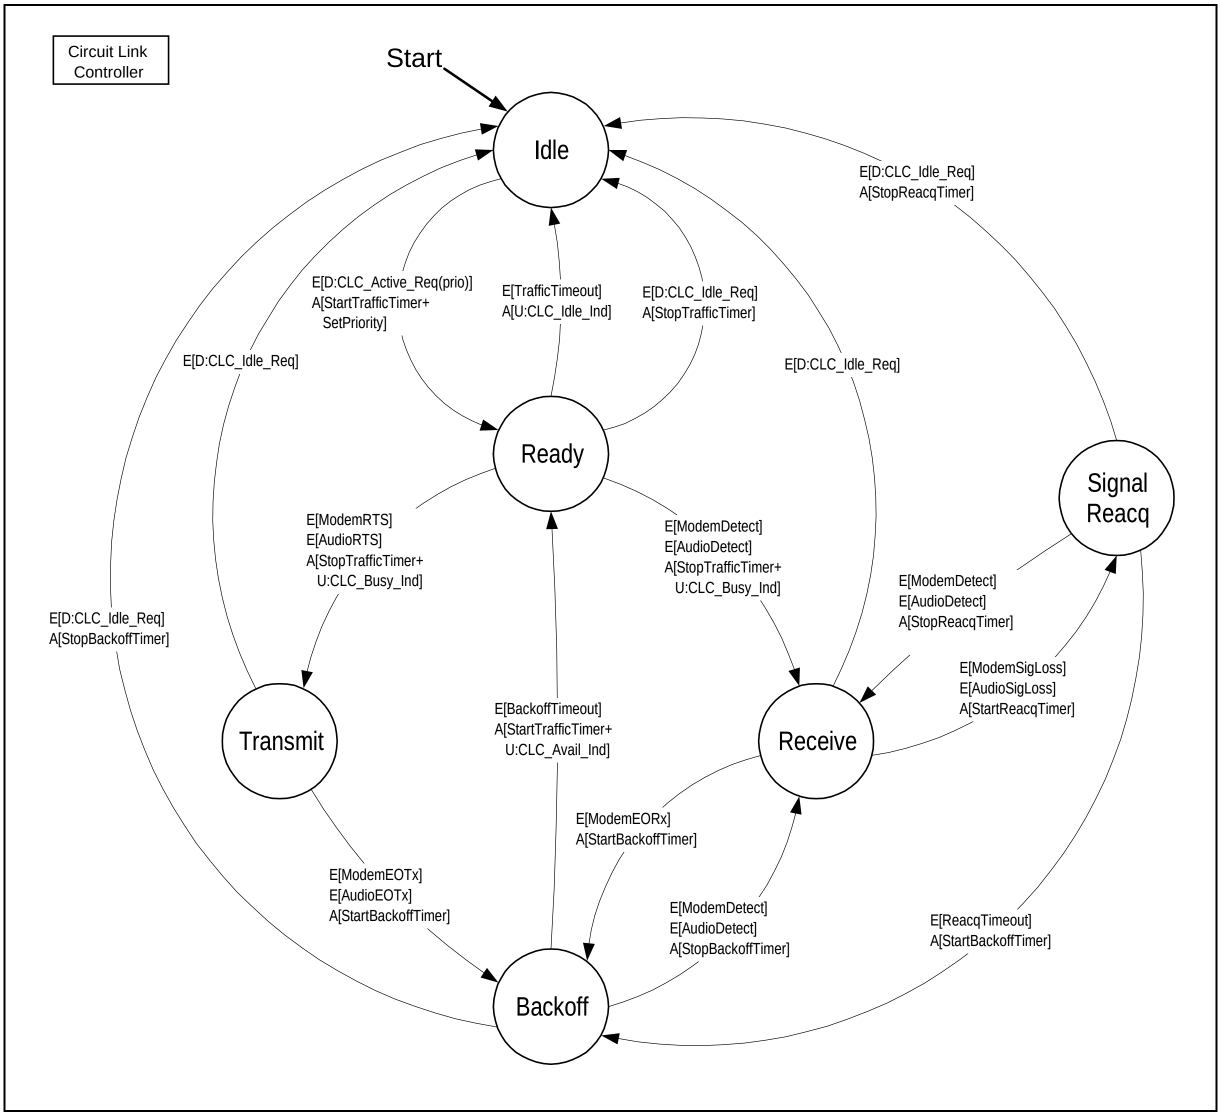

FIGURE 4.22

CLC state diagram

4.18 Simulation framework

This section describes the simulation framework used to predict the performance of 3G-HF subnetworks. Specifically, this section will discuss the modeling approach to be used and the details of physical layer modeling. In addition, waveform performance modeling and the application of the modeling framework to developing performance expectations for 3G-HF subnetworks will be discussed briefly.

4.19 Modeling approach

MIL-STD-188-141B Appendix C defines a synchronous ALE scheme for media access and two synchronous data link protocols to be used to convey data once a link has been established. Heretofore, the performance of ALE protocols and of data link protocols have been treated as separate topics. With ALE protocols, the performance metrics of concern are probability of linking given channel conditions, time to link given channel conditions, and optimality/suitability of the selected channel. With data link protocols, the performance metric of interest is typically data throughput (in bits per second) as a function of message size and channel conditions. However, for HF *subnetworks* (combining ALE, data link, and probably higher-layer protocols), the performance metrics that are ultimately of interest are message delivery rate and message delivery delay as a function of message generation rate, message size, number of network nodes, propagation conditions, and available bandwidth. The performance metrics for the ALE and data link protocols in isolation certainly provide some insight into HF subnetwork performance. However, these metrics by themselves are insufficient to predict HF subnetwork performance for subnetworks containing more than two nodes, as they do not account for the effects of arbitration failures and resulting collisions, collision avoidance mechanisms (if any), message routing, and the allocation of bandwidth between the ALE and data link protocols.

There presently exists no HF-industry-standard framework to permit evaluation and comparison of HF subnetwork designs in a simulation environment. A framework for simulation of HF subnetwork performance is required which provides:

1. Simulation of multiple network nodes (stations) in varying and dynamic network topologies.
2. Simulation of varying and dynamic network traffic loads.
3. Simulation of the media access and associated client protocols, and data link and associated client protocols within each network node.
4. Simulation of the physical layer to include:
    - HF skywave, Near Vertically Incident Skywave (NVIS), and groundwave propagation characteristics.
    - The time-varying nature of HF propagation.
    - Waveform performance in such propagation conditions.
    - The effects of on-air traffic collisions.
    - The effects of collocation and RF interference.
5. Results that are in good agreement with empirical data from on-air testing.

In the absence of an HF industry standard, the simulation framework described herein is based on the commercially-available event-driven simulation framework, OPNET[1]. This tool readily satisfies requirements 1, 2, and 3 above and is flexible enough to support the addition of an HF-specific physical layer model. The next three sections describe the means by which an HF physical layer model is included in the chosen framework.

4.20 Physical Layer Modeling

Ideally, the event-driven simulation framework would be mated with the same data flow simulation of the constituent 3G-HF waveforms and the Watterson channel model used to validate the waveform design to create an event-driven/data-flow simulation hybrid. However, this inclusion of a data flow simulation results in unacceptably long execution times when performed on desktop computers/workstations. As a result, some sort of simplified model of waveform performance is required to determine detection and demodulation outcomes for each constituent waveform. In turn, a simplified waveform performance model requires a propagation model, which is discussed in the next section.

[1] OPNET was recently used by the NETWARS Project Standards Working Group of the Joint Chiefs of Staff as the basis for a tactical communications modeling interoperability standard [11]. OPNET is a registered trademark of MIL3 Inc.

4.20.1 Physical Layer Modeling: Propagation

The propagation model must provide median SNR and propagation delay information as a function of time-of-day, frequency, node position, type of antenna, and antenna orientation. IONCAP [5] (or one of its more recent incarnations) is the logical choice to provide such information. Four options of increasing complexity exist:

1. Network nodes are fixed in space; median SNR and propagation delay between network nodes are fixed in time to produce a 'snapshot' of the diurnal variation.
2. Network nodes are fixed in space; median SNR and propagation delay between network nodes vary with time to reflect diurnal propagation variations.
3. Network nodes are mobile; median SNR and propagation delay between grid points are fixed in time. The actual median SNR and propagation delay is based on the grid points closest to the node pair.[2]
4. Network nodes are mobile; median SNR and propagation delay between grid points vary with time to reflect diurnal propagation variations. The actual median SNR and propagation delay are based on the grid points closest to the node pair.

Options 1-4 can all be used to specify median SNR and propagation delay for each node pair during the course of simulation. Each requires that a database of increasing complexity be generated prior to simulation and the application of suitable interpolation methods. Initially, to minimize the complexity of the model, network nodes will be fixed in space. To further minimize initial complexity, and since groundwave propagation is comparatively benign, we defer inclusion of a groundwave propagation model for a later time and thus assume nodes are spaced so as to ensure that groundwave propagation is not relevant. Similarly, collocation effects will be neglected initially.

At this point, we have a model for long-term (hourly) variation of node-pair SNR with time. Such long-term variations are required to properly stress 3G-HF connection management. An intermediate-term (fractions of minutes) variation in node-pair SNR as reported and characterized in [6] is applied to modulate the node-pair SNR values obtained from IONCAP. Such intermediate-term variations are required to properly stress third generation data link protocols. In addition, short-term (seconds and fractions of a second) variation of node-pair SNR with time due to fading is included, as it is of critical interest in a system employing short burst waveforms

Since the time constant of the short-term SNR variations can vary so as to be greater than, comparable to, or less than the duration of the waveforms of interest, an SNR profile as a function of time for each waveform arrival is used to predict detection and demodulation outcomes. A simple model of short-term node-pair SNR variation would periodically sample a random process having a particular distribution and autocorrelation function, and use this process to modulate a nominal node-pair SNR value as provided by the longer-term SNR variation processes. CCIR Reports 266-6[7] and 304-2[8] give examples of distributions and autocorrelation behavior measured for various propagation conditions and fading mechanisms. However, the choice of distribution (Rayleigh, Rician, Log-Normal, and others) and autocorrelation behavior is problematic, as it depends strongly on the fading mechanism(s) in effect, and there exists no generally-accepted method to select the appropriate distribution and autocorrelation function in effect between any two locations at any given time.

[2] If one node is located directly on a grid point, two-dimensional interpolation is available to reduce inaccuracies with respect to a direct calculation using the exact node positions [12].

For this reason, a slightly more complicated approach akin to that adopted in CCIR Recommendation 520-1[9] is used, where the dispersive properties of HF channels are modeled in the manner described by Watterson [10], having discrete, independent arrivals (time dispersion, referred to as differential time delay or multipath delay) each having IID Rayleigh amplitude (frequency dispersion, referred to as frequency spread or Doppler spread). CCIR Recommendation 520-1 assumes two discrete arrivals and provides a set of differential time delay and frequency spread values that are representative of common HF propagation conditions to be used in assessing and comparing the performance of HF data modems. It is proposed that this same framework and set of operating conditions is suitable for assessing and comparing the performance of HF networking schemes.

Thus each pair (permutation) of network nodes is provided with a fading process, as described previously, which is used to modulate the nominal SNR of the transmitted signal as determined from the long-term and intermediate-term node-pair SNR process. The fading process of each node pair is statistically independent from all other node-pair fading processes. Each node-pair has its own set of fading process parameters which can be chosen separately to be commensurate to node locations and time of day.

4.20.2 Physical layer modeling: collision

Having an SNR profile for each waveform arrival (or segment of a waveform arrival for longer transmissions) is necessary but not sufficient to determine detection and demodulation outcomes. A collision model is also required that accounts for the node-pair SNR profiles of both signal and interferer(s), and variations in arrival time. The collision model is as follows:

$$SNR_{Seff\,(dB)}(n) = 10\log_{10}\left(SNR_S(n) \Big/ \left(1+\sum_{i=0}^{M-1} SNR_{Ii}(n)\right)\right)$$

where SNR_{Seff} is the instantaneous *effective* SNR affecting receipt of the desired signal (accounting for the effects of interfering signals), SNR_S is the instantaneous signal to noise ratio of the desired signal (ignoring interference), SNR_{Ii} is the instantaneous SNR of the i^{th} interferer, and M is the number of interferers in existence during reception of the desired signal. Time dispersion (multipath delay) is integrated out prior to this calculation. This effective SNR profile is then used to predict waveform detection and demodulation outcomes.

Certain simplifying assumptions are employed in this collision model: the separate multipath arrivals from the various interferers are all uncorrelated, the signal multipath arrivals are uncorrelated, and the background noise, signal arrivals and interferer arrivals are all uncorrelated. These assumptions are generally valid for serial tone waveforms (on which the third generation protocols are based), assuming that all arrivals, taken pairwise, have time-of-arrival variations greater than one PSK symbol duration (approx. 0.5 ms). It is accepted that there will be occasions when these assumptions are not satisfied; however, such occasions should be rare in practice as a consequence of typical multipath delay values and variations in both clock drift and propagation delay.

4.21 Waveform Performance Modeling

Waveform performance models must be developed that determine detection and demodulation outcomes based on the effective SNR profile. This task has been performed for the BW0 and BW1 waveforms, and it is proposed that the framework presented here is sufficient to develop a performance model for the remaining third generation waveforms.

4.22 Application of simulation framework to third generation HF messaging protocols

At this time, the focus of this simulation effort is on third generation connection management and traffic management, as these are the least-explored protocols to date. The data link protocols have been prototyped and evaluated on-air and are well understood at this time. At the time of this writing, the physical layer model and waveform performance models for BW0 and BW1 have been implemented, and the implementation of connection management and traffic management protocols is underway.

Once the connection management and traffic management protocols have been implemented, they will be mated with existing models of the FS-1052 data link protocol and the MIL-STD-188-110A serial tone waveform to assess the performance of an HF subnetwork employing these protocols. It is understood that the performance of FS-1052 is not identical to that of the third generation data link protocols. However, this approach will provide timely insight into the performance of the connection management and traffic management protocols, since their performance is expected to be the limiting factor on aggregate HF subnetwork performance in many cases. Models of the third generation data link protocols can be incorporated as time allows, to more precisely characterize the aggregate performance of third generation HF subnetworks. It is expected that the simulation results obtained using the FS-1052 model will be conservative.

As stated previously, the HF subnetwork performance metrics that are ultimately of interest are message delivery rate and message delivery delay as a function of message generation rate, message size, number of network nodes, propagation conditions, and available bandwidth. The simulation framework described herein will be used to assess subnetwork performance in this manner.

4.23 Conclusion

This paper has described the third generation HF messaging system currently being proposed for Appendix C of the forthcoming U.S. MIL-STD-141B revision. The overall framework containing connection management, traffic management, data link protocol, and circuit link protocols has been described. Some performance data on the scaleable burst HF modem utilized for connection and traffic management have been presented. Additionally, performance data on the newly defined data link protocols have been presented and contrasted with the performance of U.S. Federal Standard 1052. Recognizing the difficulties in predicting the performance of an entire 3G-HF system, a framework for simulation, utilizing a combination of OPNET and statistical HF waveform models, has been presented.

The third generation HF messaging system presented in this paper promises to provide significant improvements in capability and performance relative to systems based on present standards such as MIL-STD-188-141A and FED-STD-1052. Work completed to date includes the initial design of the system, the design and simulation of some of the basic components of the system, and establishment of a framework suitable to the accurate modelling of the entire system.

Future work will implement the simulation capability and provide accurate performance comparisons between the third generation HF system and systems based on existing standards.